



Supermicro Total Solution for Lustre on OpenZFS

High performance file system based on open software defined storage accelerates innovation and reduces storage economics

Scale-Out Storage - Lustre

Lustre is the Parallel File system of choice for High Performance Computing (HPC) and large scale applications. To meet the capacity and throughput requirements of HPC workloads, Lustre has traditionally required adoption of custom proprietary storage products leading to vendor lock-in and reduced innovation. The evolution of software defined high availability platforms like OpenZFS running on Linux now make it possible to meet the performance requirements of HPC on open industry standard x86 platforms, reducing storage costs by up to 90%. Supermicro in Partnership with Intel® and some of the worlds leading HPC integrators offers a total solution for Lustre on OpenZFS with Supermicro's industry leading hardware, software and services infrastructure.

Open Industry Standard <Software/Defined> Architecture

- Tested and Validated Lustre Open solution based on workload optimized Supermicro Systems
- Reduce storage costs by up to 90% with Lustre on OpenZFS based on Supermicro open industry standard hardware and software
- Standard Configuration delivers industry leading storage density (90 Bay 4U) with fully redundant High Availability design
- Performance Configuration delivers unparalleled All-Flash NVMe performance in a 2U Design with full system (dual node) and data path (dual-path NVMe) redundancy

Best in Class System and Software Infrastructure

- Leverages the best of Supermicro's expansive Server and Storage product portfolio
- Based on hardened Intel Enterprise Edition for Lustre* Software on OpenZFS optimized for Supermicro systems increasing single node performance by 30%
- Flexible architecture providing a high performance, hot swappable and no single point of failure design
- High bandwidth utilizing 12Gb/s SAS3 and 56Gb/s FDR InfiniBand


Total Solution (Software, System and Service)

Pre-designed and pre-tested Total Solution for Lustre on OpenZFS provides customers the insight on where and how to deploy Lustre on OpenZFS reducing complexity and risk, controlling costs and accelerating solution time to market.

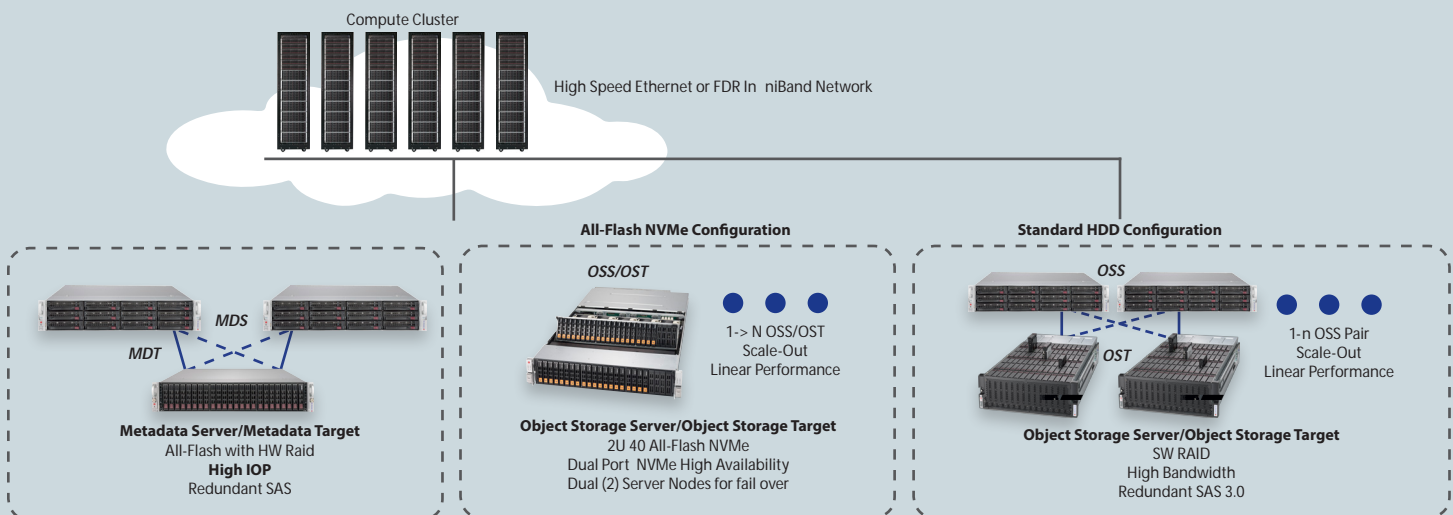
LUSTRE TOTAL SOLUTION READY HARDWARE



2U 40 Bay Dual Node Server with Dual Port All-Flash NVMe



4U 3.5" 90 Bay Dual Expander JBOD






**Object Storage Server
All-Flash NVMe**



**Object Storage Server Pod
HDD**



**Metadata and Mgmt Server Pod
(MDS/MGS)**

Description	<ul style="list-style-type: none"> • OSS: Multiple systems can be added to meet capacity requirements • Linear scaling of throughput and capacity 	<ul style="list-style-type: none"> • Multiple Pods can be added to meet capacity requirements • Linear scaling of throughput and capacity 	<ul style="list-style-type: none"> • Single MDS/MGS Pod can support hundreds of OSS Pods and thousands of clients
Capacity/Pod	30 TB raw capacity using 750GB media	720TB raw capacity using 8TB media	3.2TB raw capacity, user upgradable
Form Factor	<ul style="list-style-type: none"> • 2U Configuration • Redundant Server Nodes • Dual-Path Redundant NVMe • 40-bay Storage Target (OST) 	<ul style="list-style-type: none"> • 8U Configuration • Redundant Object Storage Servers • 1x 90-bay Object Storage Target (OST) (16x lanes per node) 	<ul style="list-style-type: none"> • 6U Configuration • Redundant Metadata servers • shared SAS3 JBOD (8x lanes per node)

Technical Specifications

MDS/MGS

CPU/Memory	Dual E5-2667 v4, 3.2 GHz, 8 core/ 128GB
Fabric	User definable supporting Ethernet, InfiniBand and Omni-Path Interconnects up to 100Gb/s
Storage Interconnect	12Gb/s (SAS 3) LSI/Avago Syncro hardware RAID controllers
Metadata Target	2U 24 2.5" Drive JBOD with shared redundant expander (8X lanes/node)

OSS/OST (HDD)

CPU/Memory	Dual E5-2667v4, 3.2 GHz, 8 core/ 128GB
Fabric	User definable supporting Ethernet, InfiniBand and Omni-Path Interconnects up to 100Gb/s
Storage Interconnect	12Gb/s (SAS 3) redundant IT mode controllers
Object Data Target	4U 90 3.5" drive JBOD(s) with shared SAS 3 redundant expander (16x lanes/node)

All-Flash NVMe Configuration Specs

OSS/OST

Nodes	2 Redundant Server Nodes in 2U
CPU/Memory	Dual E5-2667v4, 3.2 GHz, 8 core/ 128GB
Fabric	User definable supporting Ethernet, InfiniBand and Omni-Path Interconnects up to 100Gb/s
Storage Interconnect	Redundant Switched PCI-E 3.0 Architecture
Drive Support	2U x 40 All-Flash NVMe Dual Port

Management Software

- Intel Manager for Lustre
- Open Industry Standard server management (SSM)
- Remotely manage hardware and OS configuration, health & power consumption of nodes in cluster

Support & Integration

- Flexible and customizable service level agreements (SLA): 4-Hour and Next Business Day Onsite Service
- Consulting service on solution architecture design
- Engineering support on installation, configuration and testing

For more Information, call your Supermicro sales representative.

www.Supermicro.com/solutions